Bayesian Optimal Sequential Multi-Hypothesis Testing in Exponential Families

Jue Wang

arXiv:1506.08915v2 [cs.IT] 26 Jan 2017

Abstract

Bayesian sequential testing of multiple simple hypotheses is a classical sequential decision problem. However, the optimal policy is computationally intractable in general, as the posterior probability space is exponentially increasing in the number of hypotheses (the curse of dimensionality in state space). We consider a specialized problem in which observations are drawn from the same exponential family. By reconstructing the posterior probability vector from a low-dimensional diagnostic sufficient statistic, it is shown that the intrinsic dimension of the reachable posterior probability space is determined by the rank of a diagnostic matrix, which cannot exceed the number of parameters governing the exponential family, or the number of hypotheses, whichever is smaller. For univariate exponential families commonly used in practice, the probability space is of one or two dimension in most cases. Hence, the optimal policy can computed in an efficient manner. Geometric interpretation and illustrative examples are presented. Simulation studies suggest that the optimal policy can substantially outperform the existing method. The results are also extended to the sequential sampling control problem.

Index Terms

Curse of dimensionality; Dynamic programming; Exponential family; Partially observable Markov decision processes (POMDP); Sequential multi-hypothesis testing; Sampling control

I. Introduction

SEQUENTIAL multi-hypothesis testing is a generalization of standard statistical hypothesis testing to account for sequential observations and multiple alternative hypotheses. After obtaining new observations, the decision maker can stop and accept one of multiple hypotheses about the underlying statistical distribution, or wait for more observations in the hope of improving the accuracy of future decisions. The goal is to identify the true hypothesis as quickly as possible and with a desired accuracy, which can often be translated to minimizing the expected cost incurred by accepting an incorrect hypothesis and making more observations.

This problem is a classical sequential decision-making problem. It involves a trade-off between the identification accuracy and time delay, which arises in a vast array of applications including medical diagnostics [1], supervised machine learning [2], network security [3], as well as educational testing, physiological monitoring, clinical trials and military target recognition, see [4] for a comprehensive discussion.

The study of sequential hypothesis testing is originated with Wald [5], who proposed a Bayes-optimal procedure for binary simple hypotheses called the sequential probability ratio test (SPRT). In SPRT, the decision maker observes independent and identically distributed (iid.) samples of a statistical distribution, one at a time, and calculates the probability ratio reflecting how likely one distribution (hypothesis) is true when compared to the other. If this ratio strongly favors one hypothesis, then she should stop and accept that hypothesis. Otherwise, she can keep observing.

The generalization of SPRT to multiple simple hypotheses has also been considered by [6], but the Bayes-optimal policy is found to be extremely difficult to implement in practice, even for only three hypotheses. The structure of the optimal policy, on the other hand, is well understood [7]. Numerous heuristics (e.g., [8], [9], [10], [11]) and asymptotically optimal procedures (e.g., [12], [13], [14]) have been studied, except the optimal policy, as noted by a recent review [4].

Queen's University, Kingston, Ontario, Canada, K7L 3N6, jw171@queensu.ca

One can view the sequential multi-hypothesis testing problem as a special type of partially observable Markov decision process problem (POMDP) with identity transition matrix [15]. The main difficulty in generalizing the optimal policy stems from the *curse of dimensionality* in dynamic programming. In the presence of two hypotheses, it suffices to consider the posterior probability of just one hypothesis, which is a scalar. But in the face of more than two hypotheses, one must consider a posterior probability vector (also called the belief vector). The size of the posterior probability space (or belief space) increases exponentially in the number of hypotheses, making the problem notoriously difficult to solve [16].

In this paper, we first consider a specialization in which the distributions in all hypotheses come from the same exponential family. Exponential families play a central role in statistical theory and include many parametric distributions (e.g., normal, binomial, Poisson) commonly used in practice. We find that the integration of sequential hypothesis testing and exponential families gives rise to a unique property, which allows us to reconstruct the N-dimensional posterior probability vector from an r-dimensional vector, which we shall call the *diagnostic sufficient statistic*, or DSS. Here, N is the number of hypotheses and r is what we call the diagnostic intrinsic dimension, which is the rank of the *diagnostic matrix* H, an $N \times M$ matrix that include all natural parameters of the hypotheses, where M is the dimension of the exponential family. Clearly, $r = \text{rank}(H) \leq \min\{N, M\}$. That is, the diagnostic intrinsic dimension is bounded by the number of hypotheses and the dimension of exponential family. For many univariate exponential-family distributions commonly used in practice, such as normal, binomial, Poisson (see Table I), we have $M \leq 2$. This means that, for these distributions, the dynamic programming problem can be reformulated in a space with *at most two* dimension, and therefore the optimal policy can be easily found through dynamic programming, even when the number of hypotheses N is large.

It is important to note that the proposed optimal solution method is not restricted to low-dimensional exponential families. Since r is the rank of an $N \times M$ matrix, it can be much smaller than both N and M. This happens when, for example, multicollinearity is present in the natural parameters associated with different hypotheses, which gives rise to a low-rank diagnostic matrix. In this situation, the optimal solution method can be applied to high-dimensional exponential families as well.

A. Relevant Literature

Sequential multi-hypotheses testing is a classical problem in statistical decision theory, therefore the relevant literature is substantial. We can only provide a sketch in this paper, but refer the readers to [4] for a comprehensive review. In general, there are two streams of literature: Bayes-optimal policy and suboptimal policies. The stream of optimal policy has been focusing on the geometric structure of acceptance regions in the belief space. The stream of suboptimal policies focuses on practical solution procedures, which can be further divided into heuristic policies and asymptotically optimal policies. Note that "multi-hypotheses testing" in this paper refers to the "identification" problem and should not be confused with the multiple testing problem or multi-armed bandit problem.

Optimal policy: The Bayes-optimal policy was first examined by [6], who formulated the problem in the belief space and showed that the optimal acceptance region for each hypothesis is convex and contains a vertex of the probability simplex. This structure can also be represented in terms of a conditional control limit policy [17]. Sequential multi-hypothesis testing has also been integrated with change detection in [18], who also characterized the geometric properties of the acceptance regions in the belief space. However, the intrinsic complexity renders the optimal implementation impractical. In this paper, we focus on the practical solution method instead of structural results.

Suboptimal policy: Many heuristic approaches are based on the parallel implementation of multiple pairwise SPRTs. To test three hypotheses concerning the mean of normal distribution, [19] constructed two SPRTs for two different pairs of hypotheses and specified a series of heuristic decision rules. [8] extended this procedure to a general number of hypotheses. A different modification of the acceptance regions was given by [10]. Representative procedures along this line are compared by [20]. However, these methods have been developed without much consideration on optimality.

An intuitively appealing approach called the M-ary sequential probability ratio test (MSPRT) is proposed in [12]. It is a decoupled likelihood ratios test and has been shown to be asymptotically optimal when the observation costs approach zero or when the probabilities of incorrect selection approach zero [13], [21]. These limiting situations would arise where one can afford to obtain a substantial amount of information before making the final selection, or when the alternative hypotheses are easily distinguishable from each other. Asymptotically optimal solutions are also the foundation for recent developments of sequential joint detection and identification [22], decentralized sensing [23], as well as sampling control [24], [25], [15]. Note that dynamic programming is generally not involved in suboptimal policies, whereas it is almost inevitable in the search for the optimal policy.

Sequential testing with exponential family: Many heuristic policies have been developed for normal distribution [19], [8], [10], yet none claims optimality. As mentioned earlier, the Sobel-Wald procedure, as well as its extensions, are based on multiple SPRTs that are operated simultaneously, in which one must specify some coordination rules to manage potential conflicts among these parallel testings. One approach that does not involve multiple SPRTs is proposed by [26] for the testing of three hypotheses about the normal mean. However, it prohibits accepting any hypothesis at the early stage. This contradicts the optimal policy found in Section III-C1 of this paper. Detailed reviews of heuristic procedures targeting specific exponential family can be found in [27]. To our best knowledge, no optimal solution method that is scalable for multiple hypotheses about exponential families can be found in the literature.

B. Contributions and limitations

The main contribution of this paper is to show that a practical optimal solution method, in which computational complexity does not grow in response to the number of hypotheses, is possible in many practice-relevant cases. This method is based on *reconstructing* the high-dimensional belief vector using a low-dimensional diagnostic sufficient statistic.

An advantage of the proposed method is that it is compatible with any discrete prior distribution, offering great flexibility to application. However, in contrast to existing approach [6] where the acceptance thresholds are fixed and independent of the prior, our method generates a set of acceptance regions that are *non-stationary* and *prior-dependent*, which may add some complexity to the implementation. Nevertheless, numerical experiments suggest that the optimal policy can substantially outperform the popular suboptimal method when the hypotheses are difficult to differentiate and the delay penalties are high.

This paper is organized as follows. Section II describes the problem, its standard formulation and solution procedure. Section III describes the belief-vector reconstruction and the reformulation of optimality equation, illustrated with applications to open problems. Section IV compares the performance of the optimal solution with the existing suboptimal procedure. Section V extends the results to sampling control problems. The summary and discussions are given in Section VI.

II. Preliminaries

Consider a sequence of iid. observations $\{Y_1, Y_2, \ldots\}$, continuous or discrete, with probability density (or mass) function f defined on $\mathcal{Y} \subset \mathbb{R}^D$. This distribution is unknown, but there are a finite number of distinct hypotheses about it, more specifically, $H_i : f = f_i, i = 0, \ldots, N$, where $\{f_0, f_1, \ldots, f_N\}$ are known distributions and one of them is equal to f. At time k, after observing the sequence $\{Y_1, \ldots, Y_k\}$, we must choose an action among the following alternatives: stop and accept hypothesis H_i, where $i \in \mathcal{N} \triangleq \{0, \ldots, N\}$, or wait until the next period and make a new observation, Y_{k+1}. The decision process is terminated if we choose to stop. In Section V, we will consider a more general problem with multiple sampling modes.

We hope to identify the true distribution with a desirable accuracy as quickly as possible. A sequential policy $\delta = (\tau, d)$ contains a stopping time τ with respect to the historical observations, and an acceptance decision rule d taking value in the set $\mathcal{N}$. The decision process is terminated at time τ when we stop observing and, if $d = i$, we accept hypothesis H_i. We let Δ denote the set of admissible policies in

which the stopping and acceptance decisions are based on the information available at time τ. Suppose that hypothesis H_i is true (namely, the actual distribution is $f = f_i$), if we stop and accept hypothesis H_j, then a termination cost $a_{ij} \geq 0$ is incurred, where $a_{ij} = 0$ if $i = j$ (no penalty for a correct identification) and $a_{ij} \geq 0$ if $i \neq j$ (penalty for misidentification). If we wait, then an observation cost $c_i \geq 0$ is incurred per period. Before any observation is obtained, some prior belief about the true hypothesis is available. Let $0 \leq \theta_i \leq 1$ denote the prior probability that the hypothesis i is true. Clearly, $\sum_{i=0}^{N} \theta_i = 1$. Let $\boldsymbol{\theta} = (\theta_0, \theta_1, \ldots, \theta_N)$ be the prior belief vector in an N-dimensional belief space $S^N \triangleq \{\Pi = (\pi_0, \pi_1, \ldots, \pi_N) \in [0,1]^{N+1} \mid \pi_0 + \pi_1 + \cdots + \pi_N = 1\}$. Our objective is to find the Bayes-optimal policy, given the prior belief, that minimizes the total expected cost over an infinite horizon

$$R^\star = \inf_{\delta=(\tau,d)\in\Delta} \mathbb{E}_\delta \Big\{ \sum_{i=0}^{N} \tau c_i \mathbb{I}_{\{f=f_i\}} + \mathbb{I}_{\{\tau<\infty\}} \sum_{i=0}^{N}\sum_{j=0}^{N} a_{ij} \mathbb{I}_{\{f=f_i, d=j\}} | \boldsymbol{\theta} \Big\}, \tag{1}$$

where $R^\star$ is also referred to as the minimum *Bayes risk*. Given the prior belief $\boldsymbol{\theta}$, let $\Pi^k = (\pi_0^k, \pi_1^k, \ldots, \pi_N^k) \in S^N$ be the belief vector at period k, where π_j^k is the posterior probability that $f = f_j$. Clearly, $\sum_{j=0}^{N} \pi_j^k = 1$. By Bayes' rule, we have

$$\Pi^k = \mathcal{B}(\Pi^{k-1}, Y_k) \triangleq \frac{\Pi^{k-1}\mathcal{G}(Y_k)}{\Pi^{k-1}\mathcal{F}(Y_k)},$$

where $\mathcal{G}(y) \triangleq \text{diag}(f_0(y), f_1(y), \ldots, f_N(y))$ is a diagonal matrix, and $\mathcal{F}(y) \triangleq (f_0(y), f_1(y), \ldots, f_N(y))^\mathsf{T}$ is a column vector. Note that Π^k depends on the prior $\boldsymbol{\theta}$ through $\Pi_0 = \boldsymbol{\theta}$. Let $F_i(y)$ denote the distribution function corresponding to $f_i(y)$, and $\mathbb{F}(y) = (F_0(y), \ldots, F_N(y))^\mathsf{T}$.

The general dynamic programming formulation uses the belief vector $\Pi^k \in S^N$ as the state variable. The state space is the N-dimensional belief space S^N [7], with the following optimality equation:

$$V(\Pi) = \min \Big\{ V_0(\Pi), V_1(\Pi), \ldots, V_N(\Pi), V_w(\Pi) \Big\}, \tag{2}$$

$$V_j(\Pi) = \sum_{i=0}^{N} \pi_i a_{ij}, \quad j = 0, \ldots, N,$$

$$V_w(\Pi) = \sum_{i=0}^{N} \pi_i c_i + \int_{\mathcal{Y}\subset\mathbb{R}^D} V\big(\mathcal{B}(\Pi, y)\big) \Pi d\mathbb{F}(y). \tag{3}$$

One can interpret the value function $V(\Pi)$ as the minimum expected cost to go given the current belief vector $\Pi \in S^N$. $V_j(\Pi)$ is the expected cost of accepting the hypothesis j immediately. $V_w(\Pi)$ is the expected cost of deciding to wait for one more period, incur the observation cost, collect a new observation $y \in \mathcal{Y}$ at the next period, and make optimal decisions onward. This optimality equation suffers from the curses of dimensionality in both the state space and outcome space, because the belief space S^N grows exponentially in N and the integral in (3) is taken over a D-dimensional space.

It has been shown that the solution to the optimality equation is unique [7], and the optimal policy is a stationary policy that chooses the action minimizing the right-hand side of (2). Let $\Gamma_j \triangleq \{\Pi \in S^N : V(\Pi) = V_j(\Pi)\}$ be the optimal acceptance region for the hypothesis j. It is optimal to accept the hypothesis j as soon as the belief vector Π enters this region, or wait if $\Pi \notin \Gamma_j$, for all $j \in \mathcal{N}$. One can implement the optimal policy by computing all the acceptance regions and comparing the belief vector against them to make decisions. We refer to this as the *belief-vector procedure* in this paper, as illustrated in the upper panel of Figure 3. Note that this procedure has two notable features: (1) Γ_j's do not change over time. (2) Γ_j's are independent of the prior $\boldsymbol{\theta}$. Wald's SPRT and most asymptotic optimal policies such as [12], [13] also have similar features. In contrast, we will show that the optimal acceptance regions for diagnostic sufficient statistic *depend* on both time and the prior.

TABLE I
UNIVARIATE EXPONENTIAL FAMILY COMMONLY USED IN PRACTICE

Distribution	p.d.f (p.m.f)	$\boldsymbol{\alpha}$	$\boldsymbol{\eta}^{\mathsf{T}}(\boldsymbol{\alpha})$	$\boldsymbol{t}^{\mathsf{T}}(y)$	$B(\boldsymbol{\alpha})$	M
Beta	$\frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} y^{\alpha-1}(1-y)^{\beta-1}$	(α,β)	(α,β)	$(\ln y, \ln(1-y))$	$\gamma(\alpha)+\gamma(\beta)-\gamma(\alpha+\beta)$	2
Binomial	$\binom{n}{y} p^y (1-p)^{(n-y)}$	p	$\ln(\frac{p}{1-p})$	y	$-n\ln(1-p)$	1
Chi-squared	$\frac{1}{2^{\nu/2}\Gamma(\nu/2)} y^{\frac{\nu}{2}-1} e^{-\frac{y}{2}}$	ν	$\frac{\nu}{2}-1$	$\ln y$	$\ln\Gamma(\frac{\nu}{2})+\frac{\nu}{2}\ln 2$	1
Exponential	$\lambda e^{-\lambda y}$	λ	$-\lambda$	y	$-\ln\lambda$	1
Gamma	$\frac{\lambda^\alpha}{\Gamma(\alpha)} y^{\alpha-1} e^{-\lambda y}$	(α,λ)	$(\alpha,-\lambda)$	$(-\ln y, y)$	$\gamma(\alpha)-\alpha\ln\lambda$	2
Geometric	$p(1-p)^y$	p	$\ln(1-p)$	y	$-\ln p$	1
Laplace[1]	$\frac{1}{2b}\exp(-\frac{\lvert y-\mu\rvert}{b})$	b	$-1/b$	$\lvert y-\mu\rvert$	$\ln 2b$	1
Lognormal	$\frac{1}{\sqrt{2\pi\sigma^2}y}\exp\{\frac{-(\ln y-\mu)^2}{2\sigma^2}\}$	(μ,σ^2)	$(\frac{\mu}{\sigma^2},-\frac{1}{2\sigma^2})$	$(\ln y,(\ln y)^2)$	$\frac{\mu^2}{2\sigma^2}+\frac{1}{2}\ln\sigma^2$	2
Neg. binomial	$\binom{-\alpha}{y} p^\alpha (p-1)^y$	(α,p)	$\ln(1-p)$	y	$-\alpha\ln p$	1
Normal	$\frac{1}{\sqrt{2\pi\sigma^2}}\exp\{\frac{-(y-\mu)^2}{2\sigma^2}\}$	(μ,σ^2)	$(\frac{\mu}{\sigma^2},-\frac{1}{2\sigma^2})$	(y,y^2)	$\frac{\mu^2}{2\sigma^2}+\frac{1}{2}\ln\sigma^2$	2
Pareto	$\frac{\beta\alpha^\beta}{y^{(\beta+1)}}$	(α,β)	$-\beta$	$\ln y$	$-\ln\beta-\beta\ln\alpha$	1
Poisson	$\frac{\lambda^y}{y!}e^{-\lambda}$	λ	$\ln\lambda$	y	λ	1
Rayleigh	$\frac{y}{\sigma^2}\exp(-\frac{y^2}{2\sigma^2})$	σ^2	$-\frac{1}{2\sigma^2}$	y^2	$\frac{1}{2}\ln\sigma^2$	1
Weibull[2]	$\frac{\gamma}{\lambda}(\frac{y}{\lambda})^{\gamma-1} e^{-(\frac{y}{\lambda})^\gamma}$	λ	$-\frac{1}{\lambda^\gamma}$	y^γ	$\gamma\ln\lambda-\ln\gamma$	1

1. with fixed mean μ
2. with fixed shape parameter γ

III. PROPOSED SOLUTION METHOD

We first provide a brief review of distributions in the exponential family, then describe the optimal solution procedure, and finally provide illustrative examples.

Exponential families contain many important distributions widely used in practice. From now on, when we mention distribution, we refer to the distribution in an exponential family.

Definition 1 (Exponential Family): Consider a random variable Y that takes on value in some space $\mathcal{Y} \subset \mathbb{R}^D$. Let $f(y;\boldsymbol{\alpha})$ be a probability density (or mass) function parameterized by $\boldsymbol{\alpha}$ from some parameter set $\mathcal{A}$. The distribution is an M-parameter exponential family if

$$f(y;\boldsymbol{\alpha}) = h(y)\exp\left[\boldsymbol{\eta}^{\mathsf{T}}(\boldsymbol{\alpha})\boldsymbol{t}(y) - B(\boldsymbol{\alpha})\right], \quad y \in \mathcal{Y}, \tag{4}$$

for some underlying measure $h : \mathcal{Y} \to \mathbb{R}_+$, natural parameter vector $\boldsymbol{\eta} : \mathcal{A} \to \mathbb{R}^M$, natural sufficient statistic vector $\boldsymbol{t} : \mathcal{Y} \to \mathbb{R}^M$, and normalization term $B : \mathcal{A} \to \mathbb{R}$, where $\boldsymbol{\eta}^{\mathsf{T}}$ denotes the transpose of $\boldsymbol{\eta}$.

See Table I for some popular univariate distributions along with their parameters. Note that these distributions have only one or two parameters, i.e., $M \leqslant 2$. But some univariate exponential families may have $M > 2$, although they are not as common in practice as those in Table I. The results of this paper are based on the following assumption:

Assumption 1: The observations are independent and identically distributed (iid.), drawn from the same exponential family.

The iid. assumption is standard in sequential decision models. As mentioned earlier, exponential families have also been extensively studied in sequential hypotheses testing and widely used in practice. Thus, this assumption is considered standard in the literature and has important practical roots.

We will show that a high-dimensional belief vector may be reduced to a low-dimensional *diagnostic sufficient statistic* (DSS) without loss of information. Consequently, the dimension of the state space can be significantly reduced. We identify the following two opportunities for dimension reduction (as illustrated in Figure 1):

1) **Opportunity 1**: the existence of *natural sufficient statistic* in exponential families. The belief vector is often chosen as the state variable of dynamic programming because it is a sufficient statistic of the observation-control history. We note that the sufficient statistic is not necessarily unique. For example, exponential families have their own sufficient statistic: the natural sufficient statistic. The

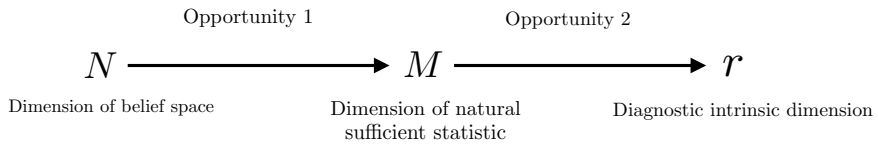

Fig. 1. Two opportunities for dimension reduction

dimension of natural sufficient statistic is denoted by M. See Table I for some examples. The first opportunity of dimension reduction is that the N-dimensional belief vector can be reconstructed from the M-dimensional natural sufficient statistic.

2) **Opportunity 2**: we can further reduce the dimensionality by exploiting the linear dependence (if any) in the *natural parameters* associated with multiple hypotheses. This opportunity is less intuitive but can be interpreted as follows: if the natural parameters of a group of hypotheses are linearly dependent, then the information that differentiates any two hypotheses can be also used to differentiate between other hypotheses in the group. This allows us to reduce the M-dimensional natural sufficient statistic further down to an r-dimensional diagnostic sufficient statistic.

We now describe the optimal solution method, which consists of two simple steps: belief vector reconstruction and the reformulation of optimality equation. Each step is detailed in the below:

A. Belief vector reconstruction

Definition 2: Define the *diagnostic matrix* as

$$H = \begin{bmatrix} \boldsymbol{\eta}^{\mathsf{T}}(\boldsymbol{\alpha}_1) - \boldsymbol{\eta}^{\mathsf{T}}(\boldsymbol{\alpha}_0) \\ \vdots \\ \boldsymbol{\eta}^{\mathsf{T}}(\boldsymbol{\alpha}_N) - \boldsymbol{\eta}^{\mathsf{T}}(\boldsymbol{\alpha}_0) \end{bmatrix},$$

where $\boldsymbol{\eta}(\boldsymbol{\alpha}_i)$ is the natural parameter vector for distribution $f_i, i \in \mathcal{N}$.

The diagnostic matrix H is an N-by-M matrix, where N is the number of hypotheses and M is the number of parameters in the exponential family. Let $r = \text{rank}(H)$ denote its rank. Clearly, $r \leq \min\{N, M\}$. Since any matrix has a rank factorization, we can always find two matrices, L and U, such that

$$H = LU,$$

where L is an N-by-r matrix of full column rank and U is an r-by-M matrix of full row rank [1]. The full-row-rank matrix U will be used to construct the diagnostic sufficient statistic:

Definition 3: Given a sequence of observations $Y_1, \ldots, Y_k$ following the distribution f with natural sufficient statistic $\boldsymbol{t}(\cdot)$, and the diagnostic matrix with rank factorization $H = LU$, we define the *diagnostic sufficient statistic* (DSS) as

$$\boldsymbol{x}^k = U \sum_{m=1}^{k} \boldsymbol{t}(Y_m).$$

Note that $\boldsymbol{x}^k$ is an r-dimensional vector, which can also be viewed as an r-dimensional projection of the cumulative sum of natural sufficient statistic. We will write it as DSS from now on. Note that it depends not only on the observations, but also the natural parameters corresponding to all hypotheses. We use $\Pi^k(\boldsymbol{Y}; \boldsymbol{\theta}) \triangleq (\pi_0^k, \ldots, \pi_N^k)$ to denote the belief vector given the iid. observations $\boldsymbol{Y} = (Y_1, \ldots, Y_k)$ and the prior belief $\boldsymbol{\theta} = (\theta_0, \ldots, \theta_N)$. The following proposition suggests that the $(N+1)$-dimensional belief vector $\Pi^k(\boldsymbol{Y}; \boldsymbol{\theta})$ can be reconstructed from the r-dimensional vector $\boldsymbol{x}^k$.

[1] Although rank factorization may not be unique, one may choose any one for our purpose.

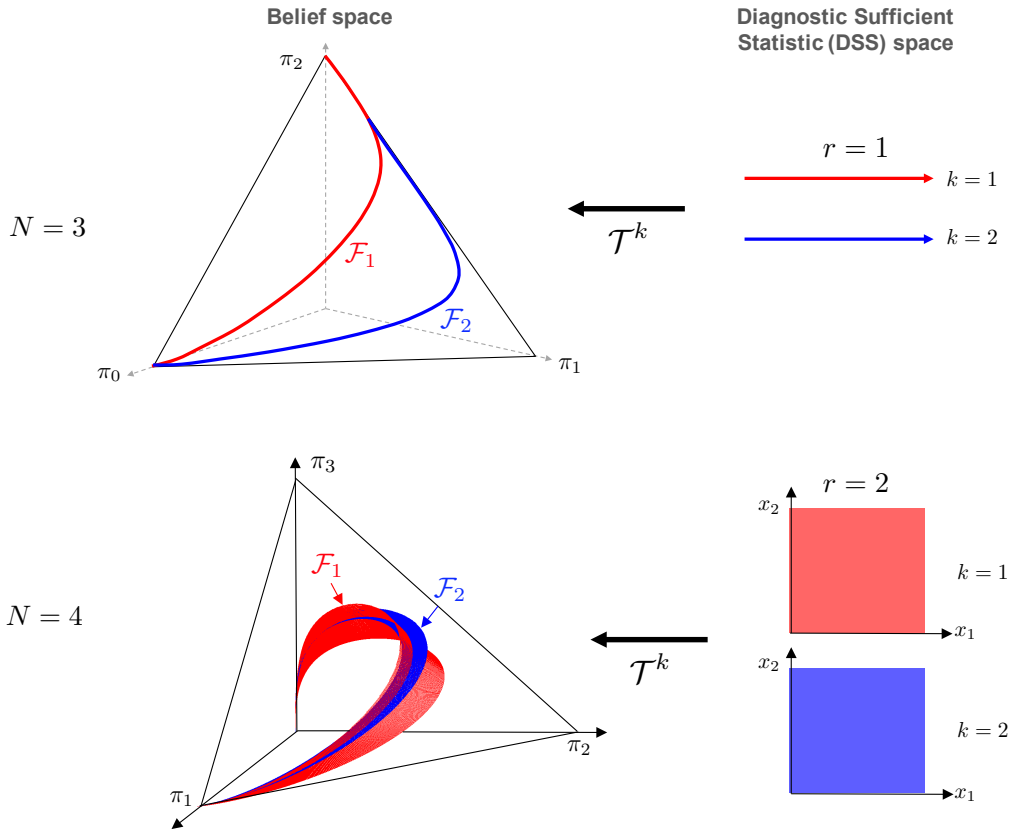

Fig. 2. Illustration of the reachable belief space.

Proposition 1 (Belief vector reconstruction): Under Assumption 1, the belief vector $\Pi^k(\boldsymbol{Y};\boldsymbol{\theta})$ can be reconstructed from the DSS, $\boldsymbol{x}^k$, through a mapping $\mathcal{T}^k : \mathbb{R}^r \to S^N$. That is, $\Pi^k(\boldsymbol{Y};\boldsymbol{\theta}) = \mathcal{T}^k(\boldsymbol{x}^k;\boldsymbol{\theta})$, where $\mathcal{T}^k \triangleq (\mathcal{T}_0^k, \ldots, \mathcal{T}_N^k)$,

$$\pi_0^k = \mathcal{T}_0^k(\boldsymbol{x}^k;\boldsymbol{\theta}) \triangleq \Big(\sum_{i=1}^N \frac{\theta_i}{\theta_0} \exp\left\{D_i^k(\boldsymbol{x}^k)\right\} + 1\Big)^{-1},$$

$$\pi_j^k = \mathcal{T}_j^k(\boldsymbol{x}^k;\boldsymbol{\theta}) \triangleq \frac{\theta_j \exp\left\{D_j^k(\boldsymbol{x}^k)\right\}}{\sum_{i=1}^N \theta_i \exp\left\{D_i^k(\boldsymbol{x}^k)\right\} + \theta_0}, \quad j = 1, \ldots, N,$$

$D_i^k(\boldsymbol{x}^k) \triangleq \boldsymbol{e}_i L \boldsymbol{x}^k - k\big[B(\boldsymbol{\alpha}_i) - B(\boldsymbol{\alpha}_0)\big]$, $\boldsymbol{e}_j$ is an N-dimensional unit row vector with 1 at the jth component.

Remark 1: Proposition 1 suggests that the belief vectors of interest can be represented by a DSS with the dimension $r = \text{rank}(H) \leqslant \{N, M\}$. Note in Table I that $M \leqslant 2$ in many distributions commonly used in practice, meaning that the belief space reachable from the prior, $\boldsymbol{\theta}$, has an intrinsic dimension of one or two, even when the number of hypotheses, N, is large. A geometric interpretation is given in Figure 2: the subspace of the belief space reachable in k periods from $\boldsymbol{\theta}$ is an r-dimensional manifold ($\mathcal{F}_k$) embedded in the N-dimensional belief space. The nonlinear transformation $\mathcal{T}^k$ "curls" the DSS space to form such manifold with the same intrinsic dimension as itself.

B. Reformulating the Optimality Equation

Now we reformulate the original optimality equation (2) in the DSS space. Unlike the belief vector, the update of the DSS does not require the direct use of Bayes' theorem; it only involves a simple summation, namely, $\boldsymbol{x}^{k+1} = \boldsymbol{x}^k + U\boldsymbol{t}(Y_{k+1})$, as implied by Definition 3. Consider a new value function $J^k(\boldsymbol{x};\boldsymbol{\theta}) \triangleq$

TABLE II
BELIEF-VECTOR VS. DSS APPROACH

	Belief-vector aprch.	DSS aprch.
Dimension of the state space	N	$r = \text{rank}(H) \leq \min\{N, M\}$
State variable	Π^k	$\boldsymbol{x}^k$
Value function	$V(\Pi^k)$	$J^k(\boldsymbol{x}^k; \boldsymbol{\theta})$
Stationary acceptance region	Yes	No
Prior-dependent acceptance region	No	Yes
Prior-dependent state	Yes	No

$V(\Pi^k)$, defined by moving the time index k out of the original value function $V(\cdot)$ and emphasizing Π^k's dependence on $\boldsymbol{\theta}$. The optimality equation involving $J^k(\boldsymbol{x}; \boldsymbol{\theta})$ is given in the following corollary:

Corollary 1: Under Assumption 1, the optimality equation (2) can be reformulated using the DSS, $\boldsymbol{x}$, as follows:

$$J^k(\boldsymbol{x}; \boldsymbol{\theta}) = \min\Big\{J_0^k(\boldsymbol{x}; \boldsymbol{\theta}), \ldots, J_N^k(\boldsymbol{x}; \boldsymbol{\theta}), J_w^k(\boldsymbol{x}; \boldsymbol{\theta})\Big\}, \tag{5}$$

$$J_j^k(\boldsymbol{x}; \boldsymbol{\theta}) = \sum_{i=0}^{N} \mathcal{T}_i^k(\boldsymbol{x}; \boldsymbol{\theta}) a_{ij}, \quad j = 0, \ldots, N,$$

$$J_w^k(\boldsymbol{x}; \boldsymbol{\theta}) = \sum_{i=0}^{N} \mathcal{T}_i^k(\boldsymbol{x}; \boldsymbol{\theta}) c_i + \int_{\mathcal{Y}} J^{k+1}\big(\boldsymbol{x} + U\boldsymbol{t}(y)\big) \sum_{i=1}^{N} \mathcal{T}_i^k(\boldsymbol{x}; \boldsymbol{\theta}) dF_i(y).$$

Remark 2: The function $J^k(\boldsymbol{x}; \boldsymbol{\theta})$ is defined on the r-dimensional state space where $\boldsymbol{x}$ resides. Since $r \leqslant 2$ in many real settings listed in Table I, we can solve the (one or two-dimensional) problem with value iteration. When the distribution is continuous, some discretization of the DSS space will be required. To perform the value iteration, one will need to use the truncation method described in [7], [12]. That is, we first consider a finite-horizon problem, solve the optimality equation by backward induction, then increase the length of the horizon to a extent such that the value function converges (the convergence is guaranteed by the dominated convergence theorem). Note also that $J^k(\boldsymbol{x}; \boldsymbol{\theta})$ explicitly depends on both time k and the prior $\boldsymbol{\theta}$ (as opposed to the original value function $V(\Pi)$ that is independent of k and $\boldsymbol{\theta}$). Accordingly, the acceptance regions become non-stationary and prior-dependent. We compare and contrast these two methods in Table II, although the decisions generated by them are the same.

C. Applications to Open Problems

We now apply the DSS-based approach to test the hypotheses concerning the normal distribution. Various suboptimal procedures have been developed for these seemingly basic problems, such as [19], [8], [10], [26], reviewed by [20]. However, no procedure has yet been known to be both optimal and scalable to a large number of hypotheses.

1) Testing the mean of normal distribution: We begin with a standard problem of testing simple hypotheses about the mean of normal distribution, assuming that the variances are known. Suppose that independent scalar observations $\boldsymbol{y} = \{Y_1, \ldots, Y_k\}$ are sequentially drawn from one of $N+1$ univariate normal distributions $f_i = N(\mu_i, \sigma^2), i = 0, \ldots, N$, differing in the mean. We are concerned with the hypotheses $H_i : Y_k \sim N(\mu_i, \sigma^2), i = 0, \ldots, N$. We have obtained the prior belief vector $\boldsymbol{\theta} = (\theta_0, \ldots, \theta_N)$ reflecting our initial knowledge. The first step is to find the diagnostic matrix

$$H = \begin{bmatrix} \boldsymbol{\eta}^{\mathsf{T}}(\boldsymbol{\alpha}_1) - \boldsymbol{\eta}^{\mathsf{T}}(\boldsymbol{\alpha}_0) \\ \vdots \\ \boldsymbol{\eta}^{\mathsf{T}}(\boldsymbol{\alpha}_N) - \boldsymbol{\eta}^{\mathsf{T}}(\boldsymbol{\alpha}_0) \end{bmatrix} = \begin{bmatrix} (\mu_1 - \mu_0)/\sigma^2, 0 \\ \vdots \\ (\mu_N - \mu_0)/\sigma^2, 0 \end{bmatrix}.$$

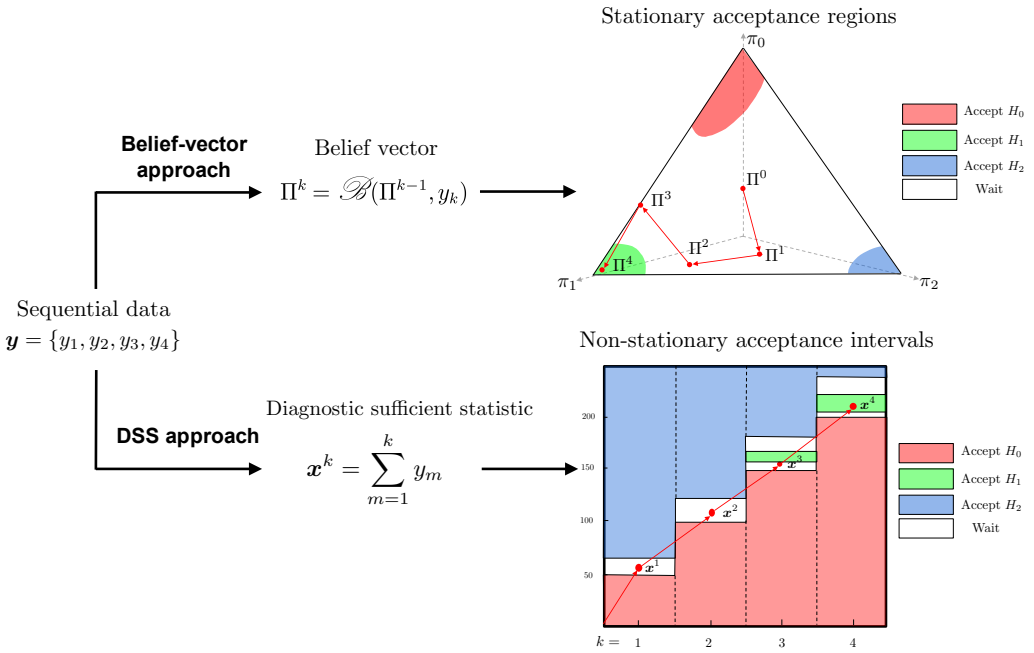

Fig. 3. Comparison of the belief-vector approach and DSS approach in Example 1.

Clearly, the rank of this matrix is one, i.e., $r = 1$. A rank factorization of H is

$$L = \big(\frac{\mu_1 - \mu_0)}{\sigma^2}, \ldots, \frac{\mu_N - \mu_0}{\sigma^2}\big)^{\mathsf{T}}, U = (1, 0).$$

The corresponding DSS is the cumulative sum of observations $\boldsymbol{x}^k = \sum_{m=1}^{k} y_m$. The transformation $\mathcal{T}^k(x; \boldsymbol{\theta}), j = 1, \ldots, N$, can be specialized as

$$\mathcal{T}_0^k(x; \boldsymbol{\theta}) = \Big(\sum_{i=1}^{N} \frac{\theta_i}{\theta_0} \exp\Big(x\frac{\mu_i - \mu_0}{\sigma^2} - k\frac{\mu_i^2 - \mu_0^2}{2\sigma^2}\Big) + 1\Big)^{-1},$$

$$\mathcal{T}_j^k(x; \boldsymbol{\theta}) = \frac{\theta_j \exp\big\{x\frac{\mu_j - \mu_0}{\sigma^2} - k\frac{\mu_j^2 - \mu_0^2}{2\sigma^2}\big\}}{\sum_{i=1}^{N} \theta_i \exp\big\{x\frac{\mu_i - \mu_0}{\sigma^2} - k\frac{\mu_i^2 - \mu_0^2}{2\sigma^2}\big\} + \theta_0}.$$

The optimality equation can be obtained by specializing (5) using $\mathcal{T}^k(x)$, defined above, the natural sufficient statistic $\boldsymbol{t}(y) = (y, y^2)^{\mathsf{T}}$, and the corresponding normal distribution function $F_i(y)$.

Example 1: Illustrative Example

Consider the example with the following parameters: $(\mu_0, \mu_1, \mu_2) = (45, 55, 60)$, $\sigma^2 = 25$, $\Pi^0 = (\theta_0, \theta_1, \theta_2) = (1/3, 1/3, 1/3)$, $(c_0, c_1, c_2) = (0.5, 0.2, 0.3)$, $a_{01} = 2, a_{02} = 5, a_{10} = 3, a_{12} = 6, a_{20} = 4, a_{21} = 7, a_{11} = a_{22} = a_{33} = 0$. The simple hypotheses to be tested are $H_0 : Y_k \sim N(45, 25), H_1 : Y_k \sim N(55, 25), H_2 : Y_k \sim N(60, 25)$. Four independent realizations of Y_k, $\{y_1, y_2, y_3, y_4\} = \{58, 52, 41, 57\}$, are sequentially generated from the distribution $N(55, 25)$.

We first describe the classical belief-vector approach. Based on the prior and sequential observation, we find the posterior belief vectors in sequence: $\Pi^1 = (0.019, 0.466, 0.515)$, $\Pi^2 = (0.013, 0.721, 0.266)$, $\Pi^3 = (0.399, 0.593, 0.008)$, $\Pi^4 = (0.039, 0.949, 0.012)$. The first three belief vectors lie in the waiting region but the fourth vector falls in the acceptance region for hypothesis H_1, as shown in Figure 3 (upper panel). Clearly, the sample path depends on the prior, but the acceptance regions do not and they remain fixed over time.

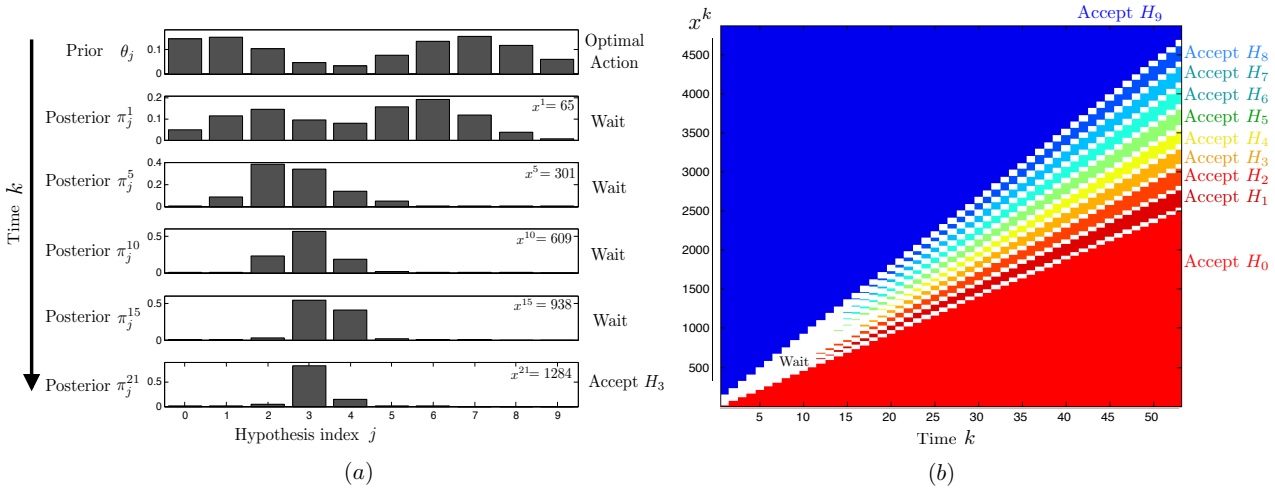

Fig. 4. Example 2: (a) belief vectors and (b) acceptance intervals on DSS.

Next, we describe the proposed solution approach. The DSS sequence is $x^1 = 58, x^2 = 110, x^3 = 151, x^4 = 208$. We compare each of them with the acceptance intervals shown in Figure 3 (lower panel) and also find that it is optimal to wait until the fourth period. The sequence of actions are, undoubtedly, the same as the belief-vector approach. But the decision process now falls in one dimension. Note that the DSS is independent of the prior, but the acceptance intervals depend on both the prior and time.

Remark 3: To implement this approach in real world, it is often desirable to use the sample average $(\sum_{m=1}^{k} y_m)/k$ as an equivalent to DSS so that the size of the state space does not increase with k. We note that the sufficient statistic $(\sum_{m=1}^{k} y_m)/k$ has been used by many heuristic methods such as Sobel-Wald procedure [19], [26]. However, the decision rules of the optimal policy draw a sharp distinction from these heuristic methods. For example, the Sobel-Wald procedure combines multiple SPRT's [19], whereas the optimal method only requires a single integrated test. The Billard-Vagholkar procedure [26] prohibits accepting any hypothesis at the first few periods, but the optimal policy allows stopping even at the first period.

Example 2: Flexible priors

The proposed method is compatible with arbitrary nonzero prior beliefs. We will illustrate such flexibility using the following example with ten hypotheses, a size that cannot be efficiently solved by the belief-vector approach.

The parameters in this example are $\mu_i = 40 + 5i$, $\sigma_i^2 = \sigma^2 = 100$, $c_i = 0.02 - 0.01i$, $a_{ij} = |i - j| + 0.5[\max(j - i, 0)]^2$, for $i = 0, \ldots, 9$. To illustrate the flexibility on prior selection, we choose a prior distribution involving a trigonometric function $\theta_i = (\sin(i) + 1.5)/16.4$ for $i = 0, \ldots, 9$, as shown in Figure 4-a. This bimodal prior implies that H_1 and H_7 are the most likely whereas H_3 and H_4 are the least likely. After collecting a series of observations, the posterior probabilities for $k = 1, 5, 10, 15, 21$ are shown in the same figure. Figure 4-b illustrates the optimal acceptance intervals on the DSS x^k for each k. These intervals are similar to those in Figure 3. Indeed, increasing the number of hypotheses no longer requires going to higher dimensions; it only adds more intervals into this chart. It is clear that the waiting intervals gradually shrink as k increases, because the uncertainty would decrease as the information accumulate.

2) Testing both mean and variance of normal distribution: When we are testing both the mean and variance, the cumulative sum of observations (y) may not be sufficient because the variance information is often better captured by the squared observations (y^2). Suppose that we sequentially observe $\boldsymbol{y} = \{Y_1, \ldots, Y_k\}$ drawn from one of $N + 1 > 2$ normal distributions $f_i = N(\mu_i, \sigma_i^2), i = 0, \ldots, N$. The goal

is to find the true distribution by testing the hypotheses $H_i : Y_k \sim N(\mu_i, \sigma_i^2), i = 0, \ldots, N$. As before, we first find the diagnostic matrix

$$H = \begin{bmatrix} \boldsymbol{\eta}^{\mathsf{T}}(\boldsymbol{\alpha}_1) - \boldsymbol{\eta}^{\mathsf{T}}(\boldsymbol{\alpha}_0) \\ \vdots \\ \boldsymbol{\eta}^{\mathsf{T}}(\boldsymbol{\alpha}_N) - \boldsymbol{\eta}^{\mathsf{T}}(\boldsymbol{\alpha}_0) \end{bmatrix} = \begin{bmatrix} \frac{\sigma_0^2\mu_1 - \sigma_1^2\mu_0}{\sigma_0^2\sigma_1^2}, \frac{\sigma_1^2 - \sigma_0^2}{2\sigma_0^2\sigma_1^2} \\ \vdots \\ \frac{\sigma_0^2\mu_N - \sigma_N^2\mu_0}{\sigma_0^2\sigma_N^2}, \frac{\sigma_N^2 - \sigma_0^2}{2\sigma_0^2\sigma_N^2} \end{bmatrix}.$$

Next, we define $\zeta_i \triangleq (\sigma_0^2\mu_i - \sigma_i^2\mu_0)/(\sigma_i^2 - \sigma_0^2)$ and examine two cases:

Case I: ζ_i's are identical.: If $\zeta_i = \zeta$ for all $i = 1, \ldots, N$, the matrix H has rank one ($r = 1$). Thus, the problem is in one dimension. This example illustrates a case where the two dimension-reduction opportunities (see Figure 1) are both present. Although the dimension of the distribution is two ($M = 2$), the intrinsic dimension of the reachable belief space is still one. That is, $r < M$. A rank factorization of H is

$$L = \big(\frac{\sigma_0^2\mu_1 - \sigma_1^2\mu_0}{\sigma_0^2\sigma_1^2}, \ldots, \frac{\sigma_0^2\mu_N - \sigma_N^2\mu_0}{\sigma_0^2\sigma_N^2}\big)^{\mathsf{T}}, U = (1, 1/2\zeta).$$

The DSS is a scalar given by

$$\boldsymbol{x}^k = U \sum_{m=1}^{k} \boldsymbol{t}(Y_m) = \sum_{m=1}^{k} (y_m + y_m^2/2\zeta).$$

Clearly, identical ζ_i's can arise when the means are identical but the variances differ ($\zeta_i = \mu_0 - \mu$), or when the means are different but the variances are identical (e.g., Example 1-2). However, it is important to note that identical ζ_i's can also appear when both mean and variance are different, for instance, when $(\mu_0, \mu_1, \mu_2, \mu_3) = (0, 1, 2, 3)$, $(\sigma_0^2, \sigma_1^2, \sigma_2^2, \sigma_3^2) = (1, 2, 3, 4)$, in which the rank of the matrix H is still one, and the corresponding DSS is still a scalar given by

$$\boldsymbol{x}^k = \sum_{m=1}^{k} (y_m + y_m^2/2).$$

But this DSS is far less intuitive than the previous cases.

Case II: ζ_i's are non-identical.: If ζ_i's are different, we have $r = 2$. A rank factorization is $L = H, U = \mathbb{I}$ (the identity matrix). In this case, the DSS is $\boldsymbol{x}^k = \sum_{m=1}^{k} \boldsymbol{t}(Y_m) = \big(\sum_{m=1}^{k} y_m, \sum_{m=1}^{k} y_m^2\big)^{\mathsf{T}}$. The transformation $\mathcal{T}^k(\boldsymbol{x}; \boldsymbol{\theta})$ can be specialized as

$$\mathcal{T}_0^k(x_1, x_2; \boldsymbol{\theta}) = \Big(\sum_{i=1}^{N} \frac{\theta_i}{\theta_0} \exp\Big\{ \frac{\sigma_0^2\mu_i - \sigma_i^2\mu_0}{\sigma_0^2\sigma_i^2} x_1 + \frac{\sigma_i^2 - \sigma_0^2}{2\sigma_0^2\sigma_i^2} x_2 - k\Big[\frac{\mu_i^2\sigma_0^2 - \mu_0^2\sigma_i^2}{2\sigma_0^2\sigma_i^2} + \ln\Big(\frac{\sigma_i}{\sigma_0}\Big)\Big]\Big\} + 1\Big)^{-1},$$

$$\mathcal{T}_j^k(x_1, x_2; \boldsymbol{\theta}) = \theta_j \exp\Big\{ \frac{\sigma_0^2\mu_j - \sigma_j^2\mu_0}{\sigma_0^2\sigma_j^2} x_1 + \frac{\sigma_j^2 - \sigma_0^2}{2\sigma_0^2\sigma_j^2} x_2 - k\big[\frac{\mu_j^2\sigma_0^2 - \mu_0^2\sigma_j^2}{2\sigma_0^2\sigma_j^2} + \ln\big(\frac{\sigma_j}{\sigma_0}\big)\big]\Big\}$$
$$\Big/ \Big\{ \sum_{i=1}^{N} \theta_i \exp\Big\{ \frac{\sigma_0^2\mu_i - \sigma_i^2\mu_0}{\sigma_0^2\sigma_i^2} x_1 + \frac{\sigma_i^2 - \sigma_0^2}{2\sigma_0^2\sigma_i^2} x_2 - k\big[\frac{\mu_i^2\sigma_0^2 - \mu_0^2\sigma_i^2}{2\sigma_0^2\sigma_i^2} + \ln\big(\frac{\sigma_i}{\sigma_0}\big)\big]\Big\} + \theta_0\Big\}, j = 1, \ldots, N.$$

We used the notation $\boldsymbol{x} = (x_1, x_2)$ in the above expressions, where $x_1 = \sum_{m=1}^{k} y_m$, $x_2 = \sum_{m=1}^{k} y_m^2$. The two-dimensional form of the finite-horizon optimality equation can be obtained by specializing (5) using the $\mathcal{T}^k(x)$ defined as above, $\boldsymbol{t}(y) = (y, y^2)^{\mathsf{T}}$, as well as the corresponding normal distribution function $F_i(y)$.

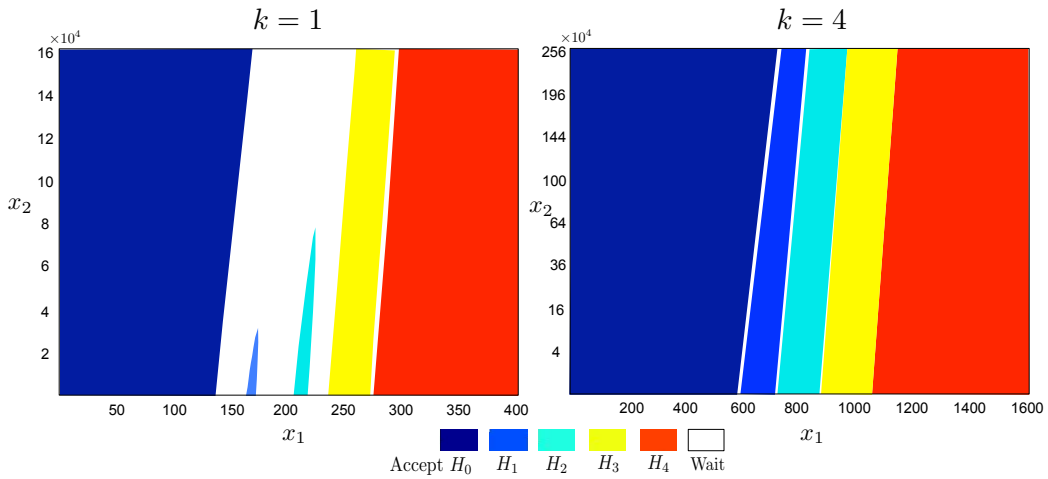

Fig. 5. The acceptance regions for Example 3.

Example 3

Consider five hypotheses about normal distribution $f_i, i = 0, \ldots, 4$, differing in both mean and variance. For the distribution f_i, the mean is given by $\mu_i = 30 + 4(i-1)^{3/2}$ and the variance follows the expression $\sigma_i^2 = 74 - i$. The prior is a uniform distribution given by $\theta_i = 1/5$. The costs are $c_i = 0.02 + 0.01i$ and $a_{ij} = |i-j|/6 + [\max(i-j, 0)]^2/12$. The DSS space has a dimension $r(H) = 2$. Examples of the acceptance regions are shown in Figure 5 for $k = 1$ and $k = 4$. The horizontal axis is the cumulative sum of observations $x_1 = \sum_{m=1}^{k} y_m$, the vertical axis is the cumulative sum of squared observations $x_2 = \sum_{m=1}^{k} y_m^2$.

IV. Comparison with MSPRT

From a practical point of view, it is important to know the magnitude of improvement that the optimal policy can provide over existing suboptimal policies. Recent developments are mainly based on asymptotically optimal policies, a good benchmark is the M-ary sequential probability ratio test (MSPRT) by [12]. This procedure is known to be asymptotically optimal as the observation costs (or identification errors) approach zero [13], [14].

Consider the case with three hypotheses about the mean of normal distribution, namely, $\mathrm{H}_i : f_i \sim N(\mu_i, \sigma^2), i = 0, 1, 2$. Suppose that the observation costs are identical, i.e., $c_i = c$ and termination costs are zero-one, namely, $a_{ij} = 1$ if $i \neq j$, $a_{ij} = 0$ if $i = j$. In this context, the MSPRT defines a series of Markov times $\tau_i = \inf\{k : \pi_i^k \geq A_i\}$, where π_i^k is the posterior probability for hypothesis i, and A_i is the corresponding constant threshold. The MSPRT stopping time is defined as the minimum Markov time $\tau = \min_i\{\tau_i\}$, and the acceptance decision rule is: $d = i$ if $\tau = \tau_i$.

We perform simulation studies to compare the performances of MSPRT with the optimal policy for different combinations of observation costs c and means μ_i. In this simulation experiment, we enumerate all combinations of MSPRT thresholds and choose the minimum cost as benchmark. The simulation is run long enough so that the width of the 95% confidence interval for estimated average cost is less than 0.001. The estimated average costs and the MSPRT's percentage of loss from optimal are shown in Table III.

We observe in Table III that the sub-optimality of MSPRT becomes larger when the observation costs c are larger, or when the differences in the mean μ_i become smaller. Such observations are consistent with asymptotic optimality. They suggest that the optimal policy is more desirable when the hypotheses are more difficult to differentiate from one another, or when we cannot afford to take many observations for fear of increasing the response delay time. Nevertheless, MSPRT gives good approximation when the

TABLE III
COMPARISON OF THE TOTAL COST BETWEEN THE OPTIMAL POLICY AND MSPRT.

$(\mu_1-\mu_0)/\sigma=0.2,\quad (\mu_2-\mu_0)/\sigma=0.4$					
c	0.5	0.4	0.3	0.2	0.1
Optimal	1.328	1.232	1.140	1.027	0.936
MSPRT	6.749	5.534	4.329	3.240	1.999
Error (%)	**408.2**	**349.2**	**279.7**	**215.5**	**113.5**
$(\mu_1-\mu_0)/\sigma=0.4,\quad (\mu_2-\mu_0)/\sigma=0.6$					
c	0.5	0.4	0.3	0.2	0.1
Optimal	1.344	1.230	1.147	1.045	0.945
MSPRT	4.181	3.315	2.821	2.140	1.452
Error (%)	**211.0**	**169.5**	**145.9**	**104.8**	**53.70**
$(\mu_1-\mu_0)/\sigma=0.8,\quad (\mu_2-\mu_0)/\sigma=1.4$					
c	0.5	0.4	0.3	0.2	0.1
Optimal	1.332	1.199	1.101	1.023	0.957
MSPRT	1.803	1.582	1.439	1.185	1.029
Error (%)	**35.39**	**31.99**	**30.84**	**15.89**	**7.497**
$(\mu_1-\mu_0)/\sigma=1.0,\quad (\mu_2-\mu_0)/\sigma=2.0$					
c	0.5	0.4	0.3	0.2	0.1
Optimal	1.280	1.161	1.120	0.980	0.976
MSPRT	1.287	1.167	1.124	0.981	0.977
Error (%)	**0.546**	**0.516**	**0.330**	**0.101**	**0.092**

hypotheses are relatively easy to differentiate and when the observation cost is low. In these situations, one may argue that MSPRT serves as a satisfactory alternative.

Incidentally, for cases in Table III, the computation time of the optimal policy ranges from 17.9 to 24.6 seconds in the MATLAB$^{\text{TM}}$ environment on a desktop computer with two 3.4 GHz Intel Core i7 processors. This is far from prohibitive for applications of hypothesis testing.

The poor performance of MSPRT seen in some situations has intuitive explanations. First, the existence of a "minimum waiting region", as shown in Figure 6-a, can increase the delay time when the optimal waiting region is actually small. Second, the MSPRT cannot fully capture the sequential nature of the problem. This can be illustrated by a simple example as follows: Consider the observation densities f_i, as shown in Figure 6-b. Let $c=0.1$ and

$$\begin{bmatrix} 0 & a_{01} & a_{02} \\ a_{10} & 0 & a_{12} \\ a_{20} & a_{21} & 0 \end{bmatrix} = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}.$$

In this setting, the misidentification cost is zero if we mistake f_1 for f_2, and vice versa; it is always one if we mistake f_0 for f_1 (or f_2), and vice versa. That is, the two hypotheses corresponding to f_1 and f_2 are *identical* in *cost*. By intuition, one may suggest grouping the two hypotheses as one hypothesis and work with the sum of probabilities $\pi_1+\pi_2$ or, equivalently, π_0. If we do so, the corresponding decision region will appear to be triangular; its boundary will be a line segment parallel to the boundary of the belief space (an example is line segment AB, as shown in Figure 6-c). As mentioned earlier, parallel boundaries of this sort are the essence of MSPRT, which sets the threshold on individual posterior probability (in this case, π_0). Under this policy, point A and B should have the same action. However, the optimal actions are actually different for the two points. To see why this is true, we should realize that f_2 is easier to differentiate from f_0 as compared with f_1, so by waiting at point A, it is more likely to obtain a new observation with a higher discriminating power (since the belief vector at point A indicates that f_2 is more likely). In other words, the expected value of new information is higher at point A; hence, one should be more "patient" and wait longer. However, this discrepancy between point A and B is lost in the sum $\pi_1+\pi_2$. If we were to make a one-shot decision during standard (non-sequential) hypothesis testing, there would be no need to dissect such details within the sum. However, this subtlety is crucial in the sequential environment and can only be captured by the optimal policy.

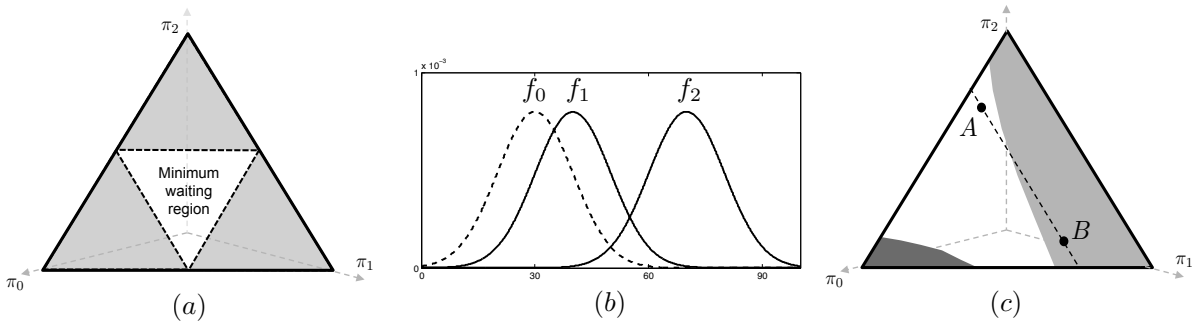

Fig. 6. The differences between MSPRT and optimal policy.

V. The Sampling Control Problem

We now extend the main results to the sequential multi-hypothesis testing problem with sampling control, in which one can adaptively choose among multiple alternative sampling modes with different diagnostic powers and costs. This subject is initiated by [24], and still remains a vibrant area of research [25], [15].

Consider a set of hypotheses denoted by $H_i, i \in \mathcal{N}$, among which only one is true. The decision maker is interested in finding the true hypothesis by conducting sequential sampling and observation. At the decision epoch k, the decision maker can either accept one of the hypotheses and terminate the decision process, or choose a sampling mode a_k from the sampling action set $\mathcal{A} = \{1, \ldots, K\}$. When the true hypothesis is H_i, the action $a \in \mathcal{A}$ generates an observation $Y_k \in \mathcal{Y} \subset \mathbb{R}^D$ with probability density (or mass) function f_i^a. We assume that the functions $f_i^a, a \in \mathcal{A}, i \in \mathcal{N}$ are known and the observations Y_k's are independent conditional on the action and hypothesis. A sequential policy $\delta = (\tau, A^\tau, d)$ contains the stopping time τ, the sequential sampling actions $A^\tau = \{a_1, \ldots, a_{\tau-1}\}$, and the acceptance decision rule $d : A^\tau \times \{Y_1, \ldots, Y_{\tau-1}\} \to \mathcal{N}$. Let $\boldsymbol{\theta} = (\theta_0, \theta_1, \ldots, \theta_N) \in S^N$ be the prior belief vector.

Suppose that f_j^a belongs to an exponential family with the natural parameter α_j^a and the natural sufficient statistic $\boldsymbol{t}_j$, namely, $f_j^a(y) = f(y; \boldsymbol{\alpha}_j^a) = h(y) \exp\left[\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_j^a)\boldsymbol{t}(y) - B(\boldsymbol{\alpha}_j^a)\right]$. Let $\boldsymbol{Y}^k = (Y_1, \ldots, Y_k)$ be the observation sequence generated by the sampling-mode sequence $A^k = (a_1, \ldots, a_k) \in \mathcal{A}^k$. Let $\Omega_a \triangleq \{m \in \{1, \ldots, k\} : a_m = a\}$ be the set of decision periods (up to time $k < \tau$) at which the sampling mode $a \in \mathcal{A}$ is used. Further, let $k_a \in \mathbb{N}$ be the cardinality of the set Ω_a, representing the total number of times of using the sampling mode $a \in \mathcal{A}$ before stopping. Clearly, $\cup_{a\in\mathcal{A}}\Omega_a^k = \{1, \ldots, k\}$, and $\sum_{a\in\mathcal{A}} k_a = k$ for $k < \tau$.

Define an $N \times (MK + K)$ diagnostic matrix as follows

$$H_s = \begin{bmatrix} \boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_1^1) - \boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0^1), & \cdots, & \boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_1^K) - \boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0^K), & B(\boldsymbol{\alpha}_0^1) - B(\boldsymbol{\alpha}_1^1), & \cdots, & B(\boldsymbol{\alpha}_0^K) - B(\boldsymbol{\alpha}_1^K) \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_N^1) - \boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0^1), & \cdots, & \boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_N^K) - \boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0^K), & B(\boldsymbol{\alpha}_0^1) - B(\boldsymbol{\alpha}_N^1), & \cdots, & B(\boldsymbol{\alpha}_0^K) - B(\boldsymbol{\alpha}_N^K) \end{bmatrix},$$

Let $r_s = \text{rank}(H_s)$ denote its rank, so $r_s \leq \min\{N, MK + K\}$. Consider a rank factorization, $H_s = L_s U_s$, where L_s is an N-by-r_s matrix of full column rank and U_s is an r_s-by-$(MK + K)$ matrix of full row rank.

Definition 4: Define the DSS as

$$\boldsymbol{x}_s^k = U_s\big(\sum_{m\in\Omega_1} \boldsymbol{t}^\mathsf{T}(Y_m), \ldots, \sum_{m\in\Omega_K} \boldsymbol{t}^\mathsf{T}(Y_m), k_1, \ldots, k_K\big)^\mathsf{T}.$$

Let $\Pi^k(\boldsymbol{Y}^k, A^k; \boldsymbol{\theta})$ be the belief vector conditional on the observation sequence $\boldsymbol{Y}^k$, sampling-mode sequences A^k, the prior belief $\boldsymbol{\theta}$ and the time index k. For brevity, we use π_i^k to denote the $(i + 1)$th component of this belief vector. The following result suggests that this belief vector can be reconstructed from the DSS, $\boldsymbol{x}_s^k$, defined above.

Proposition 2: There is a mapping $\mathcal{T}^k : \mathbb{R}^{r_s} \to S^N$, such that $\Pi^k(\boldsymbol{Y}^k, A^k; \boldsymbol{\theta}) = \mathcal{T}^k(\boldsymbol{x}_s^k; \boldsymbol{\theta})$. More specifically, $\mathcal{T}^k = (\mathcal{T}_0^k, \ldots, \mathcal{T}_N^k)$, and

$$\pi_0^k = \mathcal{T}_0^k(\boldsymbol{x}_s^k; \boldsymbol{\theta}) \triangleq \Big(\sum_{i=1}^N \frac{\theta_i}{\theta_0} \exp\big\{D_i^k(\boldsymbol{x}_s^k)\big\} + 1 \Big)^{-1},$$

$$\pi_j^k = \mathcal{T}_j^k(\boldsymbol{x}_s^k; \boldsymbol{\theta}) \triangleq \frac{\theta_j \exp\big\{D_j^k(\boldsymbol{x}_s^k)\big\}}{\sum_{i=1}^N \theta_i \exp\big\{D_i^k(\boldsymbol{x}_s^k)\big\} + \theta_0}, \quad j = 1, \ldots, N,$$

where $D_i^k(\boldsymbol{x}_s^k) = \boldsymbol{e}_i L_s \boldsymbol{x}_s^k$.

Remark 4: When the number of sampling control actions are small as compared with the number of hypotheses, it might be beneficial to use the DSS-based approach. Further, if there are many control actions but each action can cause systematic change in the observation distributions (see Example 4), then the rank of the diagnostic matrix may still be low and the DSS approach is still preferred. When there is only one sampling mode, i.e., $K = 1$, the DSS becomes $\boldsymbol{x}_s^k = U_s\big(\sum_{m\in\Omega_1} \boldsymbol{t}^{\mathsf{T}}(Y_m), k_1\big)^{\mathsf{T}} = \big((\boldsymbol{x}^k)^{\mathsf{T}}, k\big)^{\mathsf{T}}$, consisting the r-dimensional sufficient statistic $\boldsymbol{x}^k$ introduced in Definition 3 and the time index k (which accounts for the non-stationary acceptance regions). Thus, the classical problem discussed in Section III is a special case of the sampling control problem by letting $K = 1$.

Example 4: adaptive sample size

For a random variable X, suppose that the population distribution of X is normal $N(\mu, \sigma^2)$, with known variance σ^2 but unknown mean μ. To test multiple simple hypotheses about the mean, namely $H_i : N(\mu_i, \sigma^2), i = 0, \ldots, N$, the decision maker can take multiple samples at once, or one by one, before accepting a hypothesis. Generally speaking, taking multiple samples at once is not equivalent to taking the same number of samples sequentially, because the latter allows one to stop at any time, before all samples are observed. At the decision period m, she can choose the sample size $a_m \in \{1, 2, \ldots, K\}$ and observe the sample average $\bar{X} \triangleq \frac{1}{a_m}\sum_{\ell=1}^{a_m} X_\ell$. Clearly $\bar{X} \sim N(\mu, \sigma^2/a_m)$. If the hypothesis H_i is true, then $\bar{X} \sim N(\mu_i, \sigma^2/a_m)$, which implies that $f_i^{a_m} = N(\mu_i, \sigma^2/a_m), i = 0, \ldots, N$. Note that an action can cause a global change in the variances of all hypotheses. For convenience, let $\mu_0 = 0$. The diagnostic matrix becomes

$$H_s = \begin{bmatrix} \frac{\mu_1}{\sigma^2}, & 0, & 2\frac{\mu_1}{\sigma^2}, & 0, & \cdots, & K\frac{\mu_1}{\sigma^2}, & 0, & -\frac{\mu_1^2}{2\sigma^2}, & -\frac{2\mu_1^2}{2\sigma^2}, & \cdots, & -\frac{K\mu_1^2}{2\sigma^2} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{\mu_N}{\sigma^2}, & 0, & 2\frac{\mu_N}{\sigma^2}, & 0, & \cdots, & K\frac{\mu_N}{\sigma^2}, & 0, & -\frac{\mu_N^2}{2\sigma^2}, & -\frac{2\mu_N^2}{2\sigma^2}, & \cdots, & -\frac{K\mu_N^2}{2\sigma^2} \end{bmatrix},$$

whose rank is two (unless μ_i's are identical). A DSS is

$$\boldsymbol{x}_s^k = \big(\sum_{a=1}^K \big(a \sum_{m\in\Omega_a} Y_m\big), \sum_{a=1}^K a k_a\big)^{\mathsf{T}},$$

in which $\sum_{a=1}^K a k_a$ is the total number of samples taken up to the period k.

VI. Concluding Remarks

The generalization of Wald's SPRT to multiple hypotheses has been widely discussed. The structure of the optimal policy is well understood but the optimal policy itself is difficult to implement in general. We find that, for exponential families, it is possible to devise an efficient solution method, which is scalable to a large number of simple hypotheses under flexible priors in most practical cases. The method reconstructs the belief vector using the so-called diagnostic sufficient statistic and reformulate the original dynamic programming in a low-dimensional space, whose dimensionality is determined by the rank of the diagnostic matrix. The resulting control policy is distinct from the standard belief-vector approach

in the sense that the acceptance regions are non-stationary and prior-dependent. The optimal solution is particularly desirable when the alternative hypotheses are difficult to differentiate and when a quick decision has to be made.

We finally note that the diagnostic sufficient statistic is constructed using both the information from exponential family (i.e., opportunity 1) and those parameters associated with all hypotheses (i.e., opportunity 2). One needs to use the low-dimensional diagnostic sufficient statistic to reconstruct the high-dimensional belief vector and reformulate the optimality equation. This across-dimension reconstruction technique is the key to this solution.

APPENDIX A
PROOF OF PROPOSITION 1

We first show that π_0^k can be reconstructed from $\boldsymbol{x}^k$. By the definition of $\mathcal{T}_0^k$, we have

$$\begin{aligned}
\mathcal{T}_0^k(\boldsymbol{x}^k;\boldsymbol{\theta}) &= \mathcal{T}_0^k\Big(U\sum_{m=1}^k \boldsymbol{t}(Y_m);\boldsymbol{\theta}\Big) = \Big(\sum_{i=1}^N \frac{\theta_i}{\theta_0}\exp\Big\{D_i^k\big(U\sum_{m=1}^k \boldsymbol{t}(Y_m)\big)\Big\}+1\Big)^{-1}\\
&= \Big(\sum_{i=1}^N \frac{\theta_i}{\theta_0}\exp\Big\{\boldsymbol{e}_i H\sum_{m=1}^k \boldsymbol{t}(Y_m) - k\big[B(\boldsymbol{\alpha}_i)-B(\boldsymbol{\alpha}_0)\big]\Big\}+1\Big)^{-1}\\
&= \Big(\sum_{i=1}^N \frac{\theta_i}{\theta_0}\exp\Big\{\big[\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_i)-\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0)\big]\sum_{m=1}^k \boldsymbol{t}(Y_m) - k\big[B(\boldsymbol{\alpha}_i)-B(\boldsymbol{\alpha}_0)\big]\Big\}+1\Big)^{-1}\\
&= \Big(\sum_{i=1}^N \frac{\theta_i}{\theta_0}\frac{\exp\big[\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_i)\sum_{m=1}^k \boldsymbol{t}(Y_m) - kB(\boldsymbol{\alpha}_i)\big]}{\exp\big[\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0)\sum_{m=1}^k \boldsymbol{t}(Y_m) - kB(\boldsymbol{\alpha}_0)\big]}+1\Big)^{-1}\\
&= \Big(\sum_{i=1}^N \frac{\theta_i}{\theta_0}\frac{\prod_{m=1}^k\Big(h(Y_m)\exp\big[\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_i)\boldsymbol{t}(Y_m)-B(\boldsymbol{\alpha}_i)\big]\Big)}{\prod_{m=1}^k\Big(h(Y_m)\exp\big[\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0)\boldsymbol{t}(Y_m)-B(\boldsymbol{\alpha}_0)\big]\Big)}+1\Big)^{-1}\\
&= \Big(\sum_{i=1}^N \frac{\theta_i}{\theta_0}\frac{\prod_{m=1}^k f(Y_m;\boldsymbol{\alpha}_i)}{\prod_{m=1}^k f(Y_m;\boldsymbol{\alpha}_0)}+1\Big)^{-1} = \Big(\sum_{i=1}^N \frac{\pi_i^k}{\pi_0^k}+1\Big)^{-1} = \pi_0^k,
\end{aligned}$$

where the third equality follows from the definition of D_i^k and $H = LU$, the fourth equality follows from the definition of H, the fifth and sixth equalities involve some algebraic manipulation and the reconstruction of the exponential family, the seventh equality follows from the definition of exponential family and the iid. assumption, the eighth equality follows from the Baye' rule, and the last equality follows because $\sum_{i=0}^N \pi_i^k = 1$. Using the same technique, other components can also be reconstructed from $\boldsymbol{x}^k$. For $j = 1,\ldots,N$, we have

$$\begin{aligned}
\mathcal{T}_j^k(\boldsymbol{x}^k;\boldsymbol{\theta}) &= \mathcal{T}_j^k(U\sum_{m=1}^k \boldsymbol{t}(Y_m);\boldsymbol{\theta}) = \frac{\theta_j\exp\big\{D_j^k(U\sum_{m=1}^k \boldsymbol{t}(Y_m))\big\}}{\sum_{i=1}^N \theta_i\exp\big\{D_i^k(U\sum_{m=1}^k \boldsymbol{t}(Y_m))\big\}+\theta_0}\\
&= \frac{\theta_j\exp\big\{\boldsymbol{e}_j LU\sum_{m=1}^k \boldsymbol{t}(Y_m) - k\big[B(\boldsymbol{\alpha}_j)-B(\boldsymbol{\alpha}_0)\big]\big\}}{\sum_{i=1}^N \theta_i\exp\big\{\boldsymbol{e}_i LU\sum_{m=1}^k \boldsymbol{t}(Y_m) - k\big[B(\boldsymbol{\alpha}_i)-B(\boldsymbol{\alpha}_0)\big]\big\}+\theta_0}\\
&= \frac{\theta_j\exp\big\{\big[\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_j)-\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0)\big]\sum_{m=1}^k \boldsymbol{t}(Y_m) - k\big[B(\boldsymbol{\alpha}_j)-B(\boldsymbol{\alpha}_0)\big]\big\}}{\sum_{i=1}^N \theta_i\exp\big\{\big[\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_i)-\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0)\big]\sum_{m=1}^k \boldsymbol{t}(Y_m) - k\big[B(\boldsymbol{\alpha}_i)-B(\boldsymbol{\alpha}_0)\big]\big\}+\theta_0}\\
&= \frac{\frac{\theta_j\exp\left\{\sum_{m=1}^k \boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_j)\boldsymbol{t}(Y_m)-kB(\boldsymbol{\alpha}_j)\right\}}{\theta_0\exp\left\{\sum_{m=1}^k \boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0)\boldsymbol{t}(Y_m)-kB(\boldsymbol{\alpha}_0)\right\}}}{\sum_{i=1}^N \frac{\theta_i\exp\left\{\sum_{m=1}^k \boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_i)\boldsymbol{t}(Y_m)-kB(\boldsymbol{\alpha}_i)\right\}}{\theta_0\exp\left\{\sum_{m=1}^k \boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0)\boldsymbol{t}(Y_m)-kB(\boldsymbol{\alpha}_0)\right\}}+1} = \frac{\frac{\theta_j\prod_{m=1}^k h(Y_m)\exp\{\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_j)\boldsymbol{t}(Y_m)-B(\boldsymbol{\alpha}_j)\}}{\theta_0\prod_{m=1}^k h(Y_m)\exp\{\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0)\boldsymbol{t}(Y_m)-B(\boldsymbol{\alpha}_0)\}}}{\sum_{i=1}^N \frac{\theta_i\prod_{m=1}^k h(Y_m)\exp\{\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_i)\boldsymbol{t}(Y_m)-B(\boldsymbol{\alpha}_i)\}}{\theta_0\prod_{m=1}^k h(Y_m)\exp\{\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0)\boldsymbol{t}(Y_m)-B(\boldsymbol{\alpha}_0)\}}+1}
\end{aligned}$$

$$= \frac{\frac{\theta_j \prod_{m=1}^{k} f_j(Y_m)}{\theta_0 \prod_{m=1}^{k} f_0(Y_m)}}{\sum_{i=1}^{N} \frac{\theta_i \prod_{m=1}^{k} f_i(Y_m)}{\theta_0 \prod_{m=1}^{k} f_0(Y_m)} + 1} = \frac{\pi_j^k / \pi_0^k}{\sum_{i=1}^{N} \pi_i^k / \pi_0^k + 1} = \pi_j^k.$$

By now we have shown that $\mathcal{T}^k(\boldsymbol{x}^k;\boldsymbol{\theta}) = (\mathcal{T}_0^k(\boldsymbol{x}^k;\boldsymbol{\theta}),\ldots,\mathcal{T}_N^k(\boldsymbol{x}^k;\boldsymbol{\theta})) = (\pi_0^k,\ldots,\pi_N^k) = \Pi^k(\boldsymbol{Y};\boldsymbol{\theta})$, thereby completing the proof.

APPENDIX B
PROOF OF PROPOSITION 2

This proof is an extension of the proof of Proposition 1. Although there are some redundancy, we still show the full derivation for completeness. We first consider π_0^k. By the definition of $\mathcal{T}_0^k$, we have

$$\begin{aligned}
\mathcal{T}_0^k(\boldsymbol{x}_s^k;\boldsymbol{\theta}) &= \mathcal{T}_0^k\Big(U_s\big(\sum_{m\in\Omega_1}\boldsymbol{t}^\mathsf{T}(Y_m),\ldots,\sum_{m\in\Omega_K}\boldsymbol{t}^\mathsf{T}(Y_m),k_1,\ldots,k_K\big)^\mathsf{T};\boldsymbol{\theta}\Big)\\
&= \Big(\sum_{i=1}^{N}\frac{\theta_i}{\theta_0}\exp\Big\{D_i^k\big(U_s\big(\sum_{m\in\Omega_1}\boldsymbol{t}^\mathsf{T}(Y_m),\ldots,\sum_{m\in\Omega_K}\boldsymbol{t}^\mathsf{T}(Y_m),k_1,\ldots,k_K\big)^\mathsf{T}\big)\Big\}+1\Big)^{-1}\\
&= \Big(\sum_{i=1}^{N}\frac{\theta_i}{\theta_0}\exp\Big\{\boldsymbol{e}_i H_s\big(\sum_{m\in\Omega_1}\boldsymbol{t}^\mathsf{T}(Y_m),\ldots,\sum_{m\in\Omega_K}\boldsymbol{t}^\mathsf{T}(Y_m),k_1,\ldots,k_K\big)^\mathsf{T}\Big\}+1\Big)^{-1}\\
&= \Big(\sum_{i=1}^{N}\frac{\theta_i}{\theta_0}\exp\Big\{\sum_{a=1}^{K}\Big(\big[\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_i^a)-\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0^a)\big]\sum_{m\in\Omega_a}^{k}\boldsymbol{t}(Y_m)-\big[B(\boldsymbol{\alpha}_i^a)-B(\boldsymbol{\alpha}_0^a)\big]k_a\Big)\Big\}+1\Big)^{-1}\\
&= \Big(\sum_{i=1}^{N}\frac{\theta_i}{\theta_0}\exp\Big\{\sum_{m=1}^{k}\big[\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_i^{a_m})-\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0^{a_m})\big]\boldsymbol{t}(Y_m)-\sum_{m=1}^{k}\big[B(\boldsymbol{\alpha}_i^{a_m})-B(\boldsymbol{\alpha}_0^{a_m})\big]\Big\}+1\Big)^{-1}\\
&= \Big(\sum_{i=1}^{N}\frac{\theta_i}{\theta_0}\frac{\exp\big[\sum_{m=1}^{k}\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_i^{a_m})\boldsymbol{t}(Y_m)-\sum_{m=1}^{k}B(\boldsymbol{\alpha}_i^{a_m})\big]}{\exp\big[\sum_{m=1}^{k}\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0^{a_m})\boldsymbol{t}(Y_m)-\sum_{m=1}^{k}B(\boldsymbol{\alpha}_0^{a_m})\big]}+1\Big)^{-1}\\
&= \Big(\sum_{i=1}^{N}\frac{\theta_i}{\theta_0}\frac{\prod_{m=1}^{k}\Big(h(Y_m)\exp\big[\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_i^{a_m})\boldsymbol{t}(Y_m)-B(\boldsymbol{\alpha}_i^{a_m})\big]\Big)}{\prod_{m=1}^{k}\Big(h(Y_m)\exp\big[\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0^{a_m})\boldsymbol{t}(Y_m)-B(\boldsymbol{\alpha}_0^{a_m})\big]\Big)}+1\Big)^{-1}\\
&= \Big(\sum_{i=1}^{N}\frac{\theta_i}{\theta_0}\frac{\prod_{m=1}^{k}f(Y_m;\boldsymbol{\alpha}_i^{a_m})}{\prod_{m=1}^{k}f(Y_m;\boldsymbol{\alpha}_0^{a_m})}+1\Big)^{-1} = \Big(\sum_{i=1}^{N}\frac{\pi_i^k}{\pi_0^k}+1\Big)^{-1} = \pi_0^k.
\end{aligned}$$

In the fifth equality, we switch from the summation over the sampling modes to the summation over time. Applying similar arguments to other component yields:

$$\begin{aligned}
\mathcal{T}_j^k(\boldsymbol{x}_s^k;\boldsymbol{\theta}) &= \mathcal{T}_j^k(U_s\big(\sum_{m\in\Omega_1}\boldsymbol{t}^\mathsf{T}(Y_m),\ldots,\sum_{m\in\Omega_K}\boldsymbol{t}^\mathsf{T}(Y_m),k_1,\ldots,k_K\big)^\mathsf{T};\boldsymbol{\theta})\\
&= \frac{\theta_j\exp\big\{D_j^k(U_s\big(\sum_{m\in\Omega_1}\boldsymbol{t}^\mathsf{T}(Y_m),\ldots,\sum_{m\in\Omega_K}\boldsymbol{t}^\mathsf{T}(Y_m),k_1,\ldots,k_K\big)^\mathsf{T})\big\}}{\sum_{i=1}^{N}\theta_i\exp\big\{D_i^k(U_s\big(\sum_{m\in\Omega_1}\boldsymbol{t}^\mathsf{T}(Y_m),\ldots,\sum_{m\in\Omega_K}\boldsymbol{t}^\mathsf{T}(Y_m),k_1,\ldots,k_K\big)^\mathsf{T})\big\}+\theta_0}\\
&= \frac{\theta_j\exp\big\{\sum_{a=1}^{K}\Big(\big[\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_j^a)-\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0^a)\big]\sum_{m\in\Omega_a}^{k}\boldsymbol{t}(Y_m)-\big[B(\boldsymbol{\alpha}_j^a)-B(\boldsymbol{\alpha}_0^a)\big]k_a\Big)\big\}}{\sum_{i=1}^{N}\theta_i\exp\big\{\sum_{a=1}^{K}\Big(\big[\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_i^a)-\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0^a)\big]\sum_{m\in\Omega_a}^{k}\boldsymbol{t}(Y_m)-\big[B(\boldsymbol{\alpha}_i^a)-B(\boldsymbol{\alpha}_0^a)\big]k_a\Big)\big\}+\theta_0}\\
&= \frac{\theta_j\exp\big\{\sum_{m=1}^{k}\big[\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_j^{a_m})-\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0^{a_m})\big]\boldsymbol{t}(Y_m)-\sum_{m=1}^{k}\big[B(\boldsymbol{\alpha}_j^{a_m})-B(\boldsymbol{\alpha}_0^{a_m})\big]\big\}}{\sum_{i=1}^{N}\theta_i\exp\big\{\sum_{m=1}^{k}\big[\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_i^{a_m})-\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0^{a_m})\big]\boldsymbol{t}(Y_m)-\sum_{m=1}^{k}\big[B(\boldsymbol{\alpha}_i^{a_m})-B(\boldsymbol{\alpha}_0^{a_m})\big]\big\}+\theta_0}
\end{aligned}$$

$$= \frac{\frac{\theta_j \exp\left\{\sum_{m=1}^k \boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_j^{a_m})\boldsymbol{t}(Y_m) - \sum_{m=1}^k B(\boldsymbol{\alpha}_j^{a_m})\right\}}{\theta_0 \exp\left\{\sum_{m=1}^k \boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0^{a_m})\boldsymbol{t}(Y_m) - \sum_{m=1}^k B(\boldsymbol{\alpha}_0^{a_m})\right\}}}{\sum_{i=1}^N \frac{\theta_i \exp\left\{\sum_{m=1}^k \boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_i^{a_m})\boldsymbol{t}(Y_m) - \sum_{m=1}^k B(\boldsymbol{\alpha}_i^{a_m})\right\}}{\theta_0 \exp\left\{\sum_{m=1}^k \boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0^{a_m})\boldsymbol{t}(Y_m) - \sum_{m=1}^k B(\boldsymbol{\alpha}_0^{a_m})\right\}} + 1} = \frac{\frac{\theta_j \prod_{m=1}^k h(Y_m) \exp\{\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_j^{a_m})\boldsymbol{t}(Y_m) - B(\boldsymbol{\alpha}_j^{a_m})\}}{\theta_0 \prod_{m=1}^k h(Y_m) \exp\{\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0^{a_m})\boldsymbol{t}(Y_m) - B(\boldsymbol{\alpha}_0^{a_m})\}}}{\sum_{i=1}^N \frac{\theta_i \prod_{m=1}^k h(Y_m) \exp\{\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_i^{a_m})\boldsymbol{t}(Y_m) - B(\boldsymbol{\alpha}_i^{a_m})\}}{\theta_0 \prod_{m=1}^k h(Y_m) \exp\{\boldsymbol{\eta}^\mathsf{T}(\boldsymbol{\alpha}_0^{a_m})\boldsymbol{t}(Y_m) - B(\boldsymbol{\alpha}_0^{a_m})\}} + 1}$$

$$= \frac{\frac{\theta_j \prod_{m=1}^k f_j^{a_m}(Y_m)}{\theta_0 \prod_{m=1}^k f_0^{a_m}(Y_m)}}{\sum_{i=1}^N \frac{\theta_i \prod_{m=1}^k f_i^{a_m}(Y_m)}{\theta_0 \prod_{m=1}^k f_0^{a_m}(Y_m)} + 1} = \frac{\pi_j^k / \pi_0^k}{\sum_{i=1}^N \pi_i^k / \pi_0^k + 1} = \pi_j^k,$$

thereby completing the proof.

REFERENCES

[1] K. Kabasawa and S. Kaihara, "A sequential diagnostic model for medical questioning," *Med. Inform.*, vol. 6, no. 3, pp. 175–185, 1981.

[2] K. S. Fu, *Sequential Methods in Pattern Recognition and Learning*. New York: Academic Press, 1968.

[3] J. Jung, V. Paxson, A. W. Berger, and H. Balakrishnan, "Fast portscan detection using sequential hypothesis testing," in *Proceedings of IEEE Symposium on Security and Privacy*. IEEE, May 2004, pp. 211–225.

[4] A. Tartakovsky, I. Nikiforov, and M. Basseville, *Sequential Analysis: Hypothesis Testing and Changepoint Detection*, ser. Monographs on Statistics & Applied Probability. Chapman and Hall/CRC, 2014, vol. 136.

[5] A. Wald, "Sequential tests of statistical hypotheses," *Ann. Math. Statist.*, vol. 16, no. 2, pp. 117–186, 1945.

[6] A. Wald and J. Wolfowitz, "Bayes solutions of sequential decision problems," *Ann. Math. Stat.*, vol. 21, no. 1, pp. 82–99, 1950.

[7] D. Blackwell and M. A. Girschik, *Theory of Games and Statistical Decisions*, ser. Reprint of the John Wiley 1954 edition. New York: Dover, 1979.

[8] P. Armitage, "Sequential analysis with more than two alternative hypotheses, and its relation to discriminant function analysis," *J. R. Stat. Soc. Ser. B*, vol. 12, no. 1, pp. 137–144, 1950.

[9] E. Paulson, "A sequential decision procedure for choosing one of k hypotheses concerning the unknown mean of a normal distribution," *Ann. Math. Statist.*, vol. 34, no. 2, pp. 549–554, 1962.

[10] G. Simons, "A sequential three hypothesis test for determining the mean of a normal population with known variance," *Ann. Math. Statist.*, vol. 38, no. 5, pp. 1365–1375, 1967.

[11] G. Lorden, "Nearly-optimal sequential tests for finitely many parameter values," *Ann. Statist.*, vol. 5, no. 1, pp. 1–21, 1977.

[12] C. W. Baum and V. V. Veeravalli, "A sequential procedure for multi-hypothesis testing," *IEEE Trans. Inform. Theory*, vol. 40, no. 6, pp. 1994–2007, 1994.

[13] V. Dragalin, A. Tartakovsky, and V. Veeravalli, "Multihypothesis sequential probability ratio tests -Part I: Asymptotic optimality," *IEEE Trans. Inform. Theory*, vol. 45, no. 11, pp. 2448–2461, 1999.

[14] ——, "Multihypothesis sequential probability ratio tests Part II: Accurate asymptotic expansions for the expected sample size," *IEEE Trans. Inform. Theory*, vol. 46, no. 4, pp. 1366–1383, 2000.

[15] M. Naghshvar and T. Javidi, "Active sequential hypothesis testing," *Ann. Statist.*, vol. 41, no. 6, pp. 2703–2738, 2013.

[16] C. H. Papadimitriou and J. N. Tsitsiklis, "The complexity of Markov decision processes," *Math. Oper. Res.*, vol. 12, pp. 441–450, 1987.

[17] A. Tartakovsky, "Sequential testing of many simple hypotheses with independent observations," *Probl. Peredachi Inf.*, vol. 24, no. 4, pp. 53–66, 1988.

[18] S. Dayanik, C. Goulding, and H. Poor, "Bayesian sequential change diagnosis," *Math. Oper. Res.*, vol. 33, pp. 475–496, 2008.

[19] M. Sobel and A. Wald, "A sequential decision procedure for choosing one of three hypotheses concerning the unknown mean of a normal distribution," *Ann. Math. Statist.*, vol. 20, no. 4, pp. 502–522, 1949.

[20] B. Eisenberg, "Multihypothesis problems," in *Handbook of Sequential Analysis*, B. K. Ghosh and P. K. Sen, Eds. New York, NY: Marcel Dekker, 1991, pp. 229–243.

[21] V. V. Veeravalli and C. W. Baum, "Asymptotic efficiency of a sequential multihypothesis test," *IEEE Trans. Inform. Theory*, vol. 41, no. 6, pp. 1994–1997, 1995.

[22] T. L. Lai, "Sequential multiple hypothesis testing and efficient fault detection–isolation in stochastic systems," *IEEE Trans. Inform. Theory*, vol. 46, no. 2, pp. 595–608, 2000.

[23] Y. Wang and Y. Mei, "Asymptotic optimality theory for decentralized sequential multihypothesis testing problems," *IEEE Trans. Inform. Theory*, vol. 57, no. 10, pp. 7068–7083, 2011.

[24] H. Chernoff, "Sequential design of experiments," *Ann. Math. Statist.*, vol. 30, no. 3, pp. 755–770, 1959.

[25] S. Nitinawarat, G. Atia, and V. V. Veeravalli, "Controlled sensing for multihypothesis testing," *IEEE Trans. Autom. Control*, vol. 58, no. 10, pp. 2451–2464, 2013.

[26] L. Billard and M. K. Vagholkar, "A sequential procedure for testing a null hypothesis against a two-sided alternative hypothesis," *J. R. Stat. Soc. Ser. B*, vol. 31, no. 2, pp. 285–294, 1969.

[27] B. K. Ghosh and P. K. Sen, *Handbook of Sequential Analysis*. New York, NY.: Marcel Dekker, 1991.